

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:36:03

PAGE 1

REFERENCE NO: 235

This contribution was submitted to the National Science Foundation as part of the NSF CI 2030 planning activity through an NSF Request for Information, https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf17031. Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

Author Names & Affiliations

- Edward Blucher - University of Chicago
- John Carlstrom - University of Chicago
- Robert Gardner - University of Chicago
- Luca Grandi - University of Chicago
- Young-Kee Kim - University of Chicago
- David Miller - University of Chicago
- Mark Oreglia - University of Chicago
- Jim Pilcher - University of Chicago
- David Schmitz - University of Chicago
- Melvin Shochet - University of Chicago
- Scott Wakely - University of Chicago

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Experimental particle physics, astronomy, astrophysics, cosmology, distributed high-throughput computation, data-intensive science

Title of Submission

Cyberinfrastructure Enabling Measurement and Discovery at the Energy, Intensity, and Cosmic Frontiers

Abstract (maximum ~200 words).

The Physical Sciences Division at the University of Chicago has come to depend on various forms of advanced cyberinfrastructure (ACI) to meet the physics goals of its research programs in experimental particle physics, experimental nuclear physics, cosmological physics, astronomy and astrophysics. Computation, storage, high bandwidth wide-area network connectivity, and advanced services provide infrastructure and capabilities that support data aggregation and synthesis from instruments, simulations, processing, and data-intensive analysis of petascale data. A common attribute is the highly collaborative nature of the research, requiring tools and mechanisms for sharing and combining data and resources across institutions and infrastructures to meet physics goals. New, highly scalable platforms are required to meet the next generation of experiments. Materializing these platforms should drive investment in national and international ACI ecosystems which include advanced networks and common substrate layers that are not specific to the physical sciences nor implemented in institutional or domain-specific silos.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:36:03

PAGE 2

REFERENCE NO: 235

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

Research in the Physical Sciences Division at the University of Chicago examines questions ranging from quarks to the cosmos to the (as yet) unseen Dark Matter and energy in the Universe. Fundamental research in particle physics at the high energy frontier with the ATLAS detector at the CERN Large Hadron Collider explores the properties of the Higgs boson while searching for evidence of physics beyond the Standard Model, including supersymmetry and new, exotic forms of matter. Detailed properties of the cosmic microwave background (CMB) radiation is observed at the third generation of the South Pole Telescope (SPT-3G) yielding new insights into the early universe epoch, cosmic inflation and gravitational lensing. The XENON1T collaboration searches for dark matter in the Gran Sasso Laboratory in Italy using noble liquid technologies we are developing. Very energetic radiation from astronomical sources is detected using the VERITAS Cherenkov imaging telescope in Arizona to study a number of topics in astrophysics including gamma ray bursts, searches for pulsars, galaxy clusters, Dark Matter, black holes and supernova remnants. Experiments in neutrino physics are testing the fundamental nature of the neutrino itself and exploring the possibility of new physics in the neutrino sector while working to develop new detector technologies toward the next generation of large-scale neutrino observatories. For example the SNO+ Collaboration is carrying out a sensitive search for neutrinoless double beta decay, which could help to reveal the fundamental nature of the neutrino and provide clues about the origin of the matter - antimatter asymmetry in the universe.

Each of these research programs involve sensitive instruments producing prodigious volumes of data analyzed by multi-institutional and international collaborative teams while also relying on large scale detector simulations and event processing to interpret those data. As such, they have come to depend on various forms of advanced cyberinfrastructure (ACI) to aggregate, process and archive raw data streams in the shared computing environments supporting the respective collaborations. As these data-intensive research programs have computation and storage requirements far beyond the resources of any individual institution, ACI must provide the mechanisms to securely transport, manage, and provide access to data at the scale of hundreds of petabytes such that processing may occur where resources become available. These resources may come from the host laboratory, the individual research computing centers of the respective institutes, national-scale high performance computing centers provided by the NSF and DOE, and shared high-throughput clusters joined in virtual cluster environments such as provided by the Open Science Grid (OSG).

In the past ten years we have drawn lessons from contributing computing resources and infrastructure to our programs at all levels: local, regional, national and international. For example, the NSF funded ATLAS Midwest Tier-2 center, led by our group in collaboration with the University of Illinois and Indiana University, has been a leading computing resource in the Worldwide LHC Computing Grid (WLCG). In a given year, approximately 1,500 physicists from 25 countries submit more than 15 million jobs, process more than 50 petabytes of LHC data while writing several petabytes of processed output delivered back to their home institutions. These computing efforts have contributed significantly to over 600 ATLAS Collaboration publications since 2009. Similarly, the XENON1T, SPT-3G and VERITAS collaborations each require multi-institutional cyberinfrastructure to enable access to data and computational resources for all their collaborators and have successfully shared and incorporated infrastructure and methods developed by the LHC experiments and the OSG.

To achieve this has required ACI partnerships on the University of Chicago campus, IT infrastructure investments from the University, the University's Research Computing Center and grants from NSF ACI programs to expand wide area network capacity and develop science network zones. High bandwidth capacity links properly peered with regional, national, and international networking organizations has been a crucial component of our infrastructure allowing resource sharing and collaborative science. We expect the institutional ACI provisioning model will change over the next decade with the advance of new processing algorithms and the hybrid infrastructures they will require to obtain scale and sustain operations.

There are significant challenges ahead presented by long range plans underway for these programs. At CERN, the High Luminosity upgrade of the LHC (HL-LHC) is scheduled during the period of 2023-2026 with physics from 2026 and the decade beyond. The increased beam intensity and pileup (overlapping collisions) will significantly increase the data complexity and volume. Preliminary estimates indicate the needed CPU capacity will be factors of 60 beyond existing levels, and storage capacities factors of 30 above the current usage. Technological and market evolution in this time frame promises to introduce heterogeneity, making it challenging to exploit simple Moore's Law scaling without software evolution. More cores on a single chip introduces new challenges in parallelism, memory consumption and data I/O and therefore will require significant R&D in algorithm and platform development. Meanwhile, as existing software was designed 15-20 years ago there will be many software sustainability challenges.

At the Intensity Frontier, the international neutrino physics community has come together to develop the Deep Underground Neutrino Experiment (DUNE), a leading-edge experiment for neutrino science and proton decay studies. This experiment, together with the facility that will support it, the Long-Baseline Neutrino Facility (LBNF), is an internationally designed, coordinated and funded program, hosted at

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:36:03

PAGE 3

REFERENCE NO: 235

the Fermi National Accelerator Laboratory (Fermilab) in Batavia, Illinois. Among the topics to be pursued are the origins of matter and whether neutrinos are the reason the universe is made of matter rather than antimatter. By exploring the phenomenon of neutrino oscillations, DUNE seeks to revolutionize our understanding of neutrinos and their role in the universe. With the world's largest cryogenic particle detector located deep underground, DUNE can also search for signs of proton decay which could reveal a relation between the stability of matter and the grand unification force. An observation of a supernova neutrino burst would allow us to peer inside a newly-formed neutron star and potentially witness the creation of a black hole. The timescale for the needed ACI lines up well with that of the HL-LHC and similar planning and R&D is underway. While Fermilab will be the natural hub for data and computation, workflow orchestration systems are being developed to reach NSF-funded institutional ACI resources for large scale background and model simulations.

At the Cosmic Frontier, the next generation, or "Stage-4", ground-based cosmic microwave background experiment (CMB-S4) is being developed. It will consist of dedicated observatories with highly sensitive superconducting cameras operating at the South Pole, in Chile, and potentially other observatories/sites in the northern hemisphere. Together these instruments will provide a dramatic leap forward in our understanding of the fundamental nature of space and time and the evolution of the Universe. CMB-S4 will be designed to cross critical thresholds in testing inflation, determining the number and masses of the neutrinos, constraining possible new light relic particles, providing precise constraints on the nature of dark energy, and testing general relativity on large scales. However, these capabilities will require similarly dramatic increases in the scale of data handling and processing capacity, and new modes of combining datasets from diverse sources to develop comprehensive views of astrophysical phenomena.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

Workshops and planning discussions are taking place in the context of the needed "software upgrades" for the HL-LHC. Considerations for sustainable software R&D for trigger processing, detector simulation, event reconstruction, visualization, data access and management, workflow and resource management, data analysis and interpretation, data and software preservation, scalable platforms for machine learning-based analyses, security and access control, workforce development (careers, staffing and training), and new parallel frameworks to take advantage of the increasing heterogeneity of the hardware platforms. It is generally foreseen that in order to reach the scales implied by the accelerator and detector upgrades that the cyberinfrastructure itself needs to become more intelligent and capable, assuming more responsibility local to the resource. Indeed, the challenges of delivering HL-LHC scale computing in less than ten years under tightly constrained costs are significant. Intelligent, programmable cyberinfrastructure, or so-called "cyberinfrastructure as code", is seen as a key enabler to managing person-power costs while still realizing large scale aggregates of CPU and storage resources. Given that computing in high energy physics rests on the principles of high-throughput (owing to the fact that individual detector triggers are independent), distributed resources can be marshaled to meet overall resource demands if they can be managed effectively. Data storage, discovery, delivery, and caching will present the most serious challenge in the HL-LHC era. A long period of software R&D and engineering will be necessary to drive innovation and the delivery needed ACI. Such infrastructure must also be designed so that secure, reliable and robust preservation of those data is possible with minimal additional cost and person-power resources.

High capacity networks have been fundamental to the success of LHC, XENON1T, SPT-3G and VERITAS computing to date. The continued success of each of these scientific programs will inevitably necessitate more capabilities such that software defined ecosystems can be easily created and destroyed as driven by science workflow requirements. At UChicago we've come to rely on multiple 100 Gbps links to ESnet, Internet2 and international networks peered in Chicago such as LHCONE. We've aggressively used the network as a third computing resource by building federated data systems that allow remote access and caching of data hosted on our campus and at over 60 data centers worldwide. As the lower levels of the fabric gain intelligence we see a larger role for this mode of computation, allowing greater economy and flexibility in data placement and the location of the computational resources.

Researchers at our university require global connectivity. Devices, facilities, resources and collaborators are distributed everywhere. In order to accelerate our science, ACI on our campus must provide access to anyone, anywhere, at any time and at wide ranging scales. We therefore see an expanded role for services targeted at the campus edge so that scalable platforms that span multiple geographic locations can be easily built and operated, in some cases by remote expert teams overseeing centralized operations and service updates. Facilitating the use of hybrid infrastructures -- composed of university computing resources, national-scale HPC resources, resources shared with other campuses, and commercially available cloud resources -- will be essential in order to accommodate peak demands and deadline-

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:36:03

PAGE 4

REFERENCE NO: 235

scheduling.

For example, XENON1T today requires a uniform interface to allocated resources from the University's Research Computing Center's Midway cluster, the "backfill" capacity of the ATLAS Midwest Tier-2 cluster, the shared opportunistic cycles of the Open Science Grid, the dedicated allocations on the European data grid, and XSEDE resources such as the Comet machine at the San Diego Supercomputer Center. Flexible and easily provisioned ACI is needed to support these kinds of hybrid environments, including access to research and commercial cloud providers to overflow for peak demand or when the economics are justified.

Automation and orchestration will play critical roles in managing a new breed of distributed processing, machine learning analytics and data delivery platforms. Our legacy computing models, built as a hierarchy of data centers with increasing capacity and specific roles, has already evolved into a heterogeneous collection of sites performing tasks originally constrained to specific tiers. Advanced services for data transfer, data management, data streaming and data caching are being pushed from the data center to the wide area network creating ACI which will demand leveraging industry technology trends to achieve affordable and sustainable solutions. Mobility of capabilities across geographic regions and administrative boundaries is key theme moving forward which requires development of ubiquitous cyberinfrastructure underlayment and services.

As in existing computing models and infrastructure, opportunistically leveraging non-dedicated resources provided by third-parties will be necessary to fulfill resource requirements. These resources are expected to become increasingly diverse as compared to the relative uniformity in x86 CPU architecture, Linux operating systems in clustered environments we've enjoyed since 2000. Incorporating GPUs and other processor-accelerator technologies, advanced policy driven scheduling systems, container execution environments, and edge caching services into global workload systems for collaborations will require new kinds of information and discovery services. Automated, machine learning-based decision engines are seen as a viable path towards dynamically optimizing globally distributed resources. Comprehensive aggregations of metadata associated with ACI systems are necessary to provide sustainable operations, robustness, and fault-tolerance. Supporting collaborative teams and the systems and services they rely on at all scales will remain a challenge.

ACI supporting data curation and software (and operating system platform) preservation associated with scientific results is necessary for data and knowledge stewardship. Analysis reproducibility and extension when new datasets become available, and from one generation of students to the next, are capabilities we see as essential for advancing the research goals within our respective domains. New methods which capture the provenance of datasets, software dependencies, run time environments, and associated orchestrated workflows need to be adapted to our group, departmental, and university-wide systems.

Fundamental to these goals is having our University participate in a global network architectures and a global peerings of advanced services. Realizing this vision will require continuous investments in ACI on campus and with our collaborative partner organizations. ACI in this context means leveraging the full spectrum of equipment, services, and human resources required to deliver these capabilities.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

While our perspectives are primarily informed by provisioning and operating ACI needed for data-intensive computation at the frontiers of the physical sciences by medium to large collaborations, communities like the OSG consortium offers a context to empower a diversity of research groups having similar data-intensive workloads suitable for distributed high-throughput computing methods and technology. OSG Connect, part of the cyber-ecosystem which enables hundreds of individual researchers access the shared infrastructure of the OSG, was developed on top of automated cyberinfrastructure deployed and operated for the ATLAS Midwest Tier-2 center at the University of Chicago. Leveraging job routing software from HTCondor, federated identity management from CI-Logon and InCommon, group authorization and reliable data transfer software from Globus, and the science DMZ engineered by the University of Chicago IT Services organization, the platform provides a launch point into campus clusters of the OSG and other resource targets. Research groups from more than 60 institutions representing over 30 scientific domains have used this extension of the LHC computing infrastructure for simulations, analysis, modeling and interpretation of their data. The domains include medical science (neuroimaging of brain image scans), bioinformatics, evolutionary biology, genomics, zoology, chemistry, mathematics and electrical engineering among others. In aggregate the projects consume upwards of 60 million CPU-hours in a year. Furthermore we have enabled a diversity of academic institutions gaining access to shared, national-scale cyberinfrastructure ranging from high schools to R1 universities and national laboratories. Federating capabilities and resource sharing concepts were built into the project from the beginning.

Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 12:36:03

PAGE 5

REFERENCE NO: 235

Overall our view is that ACI investments made at resource providing institutions should be engineered and driven by policies that promote multi-institutional sharing realized by integration into the national ACI fabric from the beginning rather than as an afterthought. Provincial islands of closed, institutional, domain-silo'd HPC resources need to open up and couple with evolved science DMZ services (beyond the canonical data transfer and perfSONAR nodes) and present more flexible resource scheduling services such as for containerized job payloads and streaming content delivery networks. Open science requires open, outwardly-facing infrastructure configured with sensible security, authorization and access policies. Indeed, investments by the NSF over the past several years to strengthen institutional research computing, wide-area network connectivity and engineering best-practices have laid the groundwork for new forms of ACI that have the potential to more directly connect research teams to shared computational and data resources and to each other. We have seen how these approaches have transformed Big Science, democratizing access and capabilities. Institutional and program funded initiatives can simultaneously meet institutional and stakeholder goals and accelerate science across all scales.

Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>)."
-